

A Simple PTAS for the Dual Bin Packing Problem and Advice Complexity of Its Online Version*

Allan Borodin¹, Denis Pankratov², and Amirali Salehi-Abari³

1 University of Toronto, Toronto, Canada
bor@cs.toronto.edu

2 University of Toronto, Toronto, Canada
denisp@cs.toronto.edu

3 Faculty of Business and IT, University of Ontario Institute of Technology,
Oshawa, Canada
abari@uoit.ca

Abstract

Recently, Renault (2016) studied the dual bin packing problem in the per-request advice model of online algorithms. He showed that given $O(1/\epsilon)$ advice bits for each input item allows approximating the dual bin packing problem online to within a factor of $1 + \epsilon$. Renault asked about the advice complexity of dual bin packing in the tape-advice model of online algorithms. We make progress on this question. Let s be the maximum bit size of an input item weight. We present a conceptually simple online algorithm that with total advice $O\left(\frac{s + \log n}{\epsilon^2}\right)$ approximates the dual bin packing to within a $1 + \epsilon$ factor. To this end, we describe and analyze a simple offline PTAS for the dual bin packing problem. Although a PTAS for a more general problem was known prior to our work (Kellerer 1999, Chekuri and Khanna 2006), our PTAS is arguably simpler to state and analyze. As a result, we could easily adapt our PTAS to obtain the advice-complexity result.

We also consider whether the dependence on s is necessary in our algorithm. We show that if s is unrestricted then for small enough $\epsilon > 0$ obtaining a $1 + \epsilon$ approximation to the dual bin packing requires $\Omega_\epsilon(n)$ bits of advice. To establish this lower bound we analyze an online reduction that preserves the advice complexity and approximation ratio from the binary separation problem due to Boyar et al. (2016). We define two natural advice complexity classes that capture the distinction similar to the Turing machine world distinction between pseudo polynomial time algorithms and polynomial time algorithms. Our results on the dual bin packing problem imply the separation of the two classes in the advice complexity world.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases dual bin packing, PTAS, tape-advice complexity

Digital Object Identifier 10.4230/OASICS.SOSA.2018.8

1 Introduction

Given a sequence of items of weights w_1, \dots, w_n and m bins of unit capacity, the dual bin packing problem asks for the maximum number of items that can be packed into the bins without exceeding the capacity of any bin.¹ The search version of this problem is to find a good packing. In the online version of this problem, the items are presented one at a time in

* Research is supported by NSERC.

¹ This terminology is somewhat unfortunate, because the dual bin packing problem is not the dual to the natural integer programming formulation of the bin packing problem. For some early results on the latter see [1, 10].



some adversarial order and the algorithm needs to make an irrevocable decision into which (if any) bin to pack the current item. The dual bin packing problem has a substantial history in both the offline and online settings starting with Coffman et al. [14]. The performance of the online algorithm is measured by its competitive ratio; that is, the worst-case ratio between the value of an offline optimal solution and the value of the solution obtained by the algorithm. It is known that the online dual packing problem does not admit a constant competitive ratio even for randomized algorithms [8, 11]. The assumption that the online algorithm does not see the future at all is quite restrictive and in many cases impractical. It is often the case that some information about the input sequence is known in advance, e.g., its length, the largest weight of an item, etc. An information-theoretic way of capturing this side knowledge is given by the *tape-advice model* [3]. In this model, an all powerful oracle that sees the entire input sequence creates a short advice string. The algorithm uses the advice string in processing the online nodes. The main object of interest here is the tradeoff between the size of advice and the competitive ratio of an online algorithm. Often, a short advice string results in a dramatic improvement of the best competitive ratio that is achievable by an online algorithm. Of course, a short advice string can be computationally difficult to obtain since the oracle is allowed unlimited power.

A related advice model is the *per-request advice model* [13]. In this model, prior to seeing the i th input item, the algorithm receives the i th advice string. Unlike the tape-advice model, the overall length of advice is always lower bounded by n in this model. Both of these advice models have recently received considerable attention in the research community (see Boyer et al [5] for an extensive survey on this topic). Recently, Renault [16] studied the dual bin packing problem in the per-request advice model. He designed an algorithm that with 1 bit of advice per request achieves a $3/2$ competitive ratio. He also showed that with $O(1/\epsilon)$ bits of advice per request it is possible to achieve a $1 + \epsilon$ competitive ratio.² In [16] Renault explicitly asked, as an open problem, to analyze the advice complexity of the dual bin packing problem in the tape advice model. In this paper, we make progress on the advice complexity needed for achieving a $(1 + \epsilon)$ competitive ratio for the online dual bin packing problem. Specifically, let s be the maximum bit size of a weight of an input item. In particular, the overall input size is $O(ns)$ bits. We present an online algorithm that with $O(\frac{s + \log n}{\epsilon^2})$ bits of advice achieves a $(1 + \epsilon)$ competitive ratio for the dual bin packing problem. Note that it is trivial to achieve optimality with $n \log_2 m$ advice bits by specifying for each input item into which bin it should be placed. When stated in the tape advice model, Renault's bound for a $(1 + \epsilon)$ competitive ratio is $\Theta(n/\epsilon)$. Our advice bound for achieving a $(1 + \epsilon)$ approximation is exponentially smaller for the regime of constant ϵ and $s = O(\log n)$. When the n item weights have $s = n$ bits of precision, we show that the dependence on s is necessary by exhibiting an $\Omega_\epsilon(n)$ lower bound on the advice necessary to achieve a $(1 + \epsilon)$ approximation.

Our main result heavily relies on a simple polynomial time approximation scheme (PTAS) for the dual bin packing problem, which constitutes the technical core of this paper. Dual bin packing is a special case of the multiple knapsack problem (MKP). In the MKP, each of the n items is described by its weight (number in $(0, 1]$) and its value (an integer). There are m knapsacks each with their own capacity. The goal is to pack a subset of items such that

² Renault states the approximation as $1/(1 - \epsilon)$ whereas we will use $(1 + \epsilon)$ which is justified since $1/(1 - \epsilon) \leq 1 + \epsilon$ for all $\epsilon \leq 2/3$. Also without loss of generality, we will sometimes say that the approximation is $1 + \Theta(\epsilon)$ since our advice bounds are asymptotic and we can replace ϵ by ϵ/c for some suitable c .

all items fit into the knapsacks without violating weight constraints and the total value of packed items is as large as possible. In the uniform MKP, capacities of the bins are equal, and, are taken to be 1 without loss of generality. Thus, the dual bin packing problem can be seen as the uniform MKP with all values being 1. It is known [9] that the dual bin packing, and consequently the MKP, is strongly NP-hard even for $m = 2$, which effectively rules out an FPTAS for these problems. This is in contrast to the standard knapsack problem and the makespan problem for a fixed number of machines where FPTAS are possible. Significant progress in the study of the MKP was made by Kellerer [15] who showed that the uniform MKP admits a PTAS. Subsequently, Chandra and Khanna [9] gave a PTAS for the general MKP. Clearly, these results also give PTAS algorithms for the dual bin packing problem. However, the PTAS algorithms provided by Kellerer, and Chekuri and Khanna, are relatively complicated algorithms with a technically detailed analysis of correctness. Our goal is to provide a simple online advice algorithm for the dual bin packing problem based on a simpler PTAS for the dual bin packing problem. Thus, as a first step, we provide a simpler PTAS and analysis for the case of the dual bin packing problem. In the second step, we use the simplified PTAS to derive our result for the tape advice-complexity of the online dual bin packing. One of the key steps in our PTAS is a dynamic programming algorithm for the dual bin packing problem with few distinct weights instead of the IP solver as in Kellerer's PTAS or the LP solver as in Chekuri and Khanna. This dynamic programming algorithm is essentially the same as the one used in the solution of the makespan problem with a bounded number of different processing times. Our PTAS and its analysis are self-contained and easy to follow. Our work highlights one of the important aspects of simple algorithms, namely, they are usually easier to modify and adapt to other problems and situations. In particular, we are able to easily adapt our simple PTAS to the setting of online tape-advice algorithms.

2 Preliminaries

The dual bin packing instance is specified by a sequence of n item weights w_1, w_2, \dots, w_n and $m \in \mathbb{N}$ bins, where $w_i \in (0, 1]$. The goal is to pack a largest subset of items into m bins such that for each bin the total weight of items placed in that bin is at most 1. The problem can be specified as an integer program as follows (notation $[n]$ stands for $\{1, \dots, n\}$):

$$\begin{aligned}
 & \max. && \sum_{i=1}^n \sum_{j=1}^m x_{ij} \\
 \text{subj. to} &&& \sum_{i=1}^n x_{ij} w_i \leq 1 && \text{for all } j \in [m] \\
 &&& \sum_{j=1}^m x_{ij} \leq 1 && \text{for all } i \in [n] \\
 &&& x_{ij} \in \{0, 1\} && \text{for all } i \in [n], j \in [m]
 \end{aligned}$$

The online First Fit algorithm FF constructs a solution by processing items in the given order w_1, w_2, \dots, w_n and placing a given item into the first bin into which it fits. First Fit Increasing algorithm FFI first orders the items by increasing weight. Let $\sigma : [n] \rightarrow [n]$ be the corresponding permutation. Then FFI runs FF on the items in the order given by σ , i.e., $w_{\sigma(1)} \leq w_{\sigma(2)} \leq \dots \leq w_{\sigma(n)}$. It is easy to see that FF has an unbounded approximation (i.e.,

competitive) ratio whereas Coffman et al [14] show that FFI has a $4/3$ approximation ratio for the dual packing problem. Note that the weights only enter the above integer program as constraints and are not part of the objective function. Thus, it is easy to see that the items in an optimal solution are, without loss of generality, a prefix of $w_{\sigma(1)}, w_{\sigma(2)}, \dots, w_{\sigma(n)}$ packed into appropriate bins.

Throughout this paper we shall always write n to mean the number of input items, m the number of bins, and s the maximum bit size of an input item (s can be thought of as the “word size” of a computer, on which the given input sequence should be processed).

An online algorithm ALG is said to achieve a *competitive ratio* c for a maximization problem if there exists a constant α such that for all input sequences I we have $\text{OPT}(I) \leq c\text{ALG}(I) + \alpha$, where $\text{ALG}(I)$ is the value of the objective that the algorithm achieves on I and $\text{OPT}(I)$ is the value achieved by an offline optimal solution. If $\alpha \leq 0$, we say that ALG achieves a *strict* competitive ratio c .

3 A Simple PTAS for the Dual Bin Packing Problem

Fix $\epsilon > 0$. In what follows, for simplicity we shall assume that ϵ is “nice”, i.e., $1/\epsilon$ is an integer, $m\epsilon$ is an integer, etc. We note that an arbitrary small nice ϵ can always be found. Let $S = \{i \mid w_i \leq \epsilon\}$ be the set of small input items, and let $L = \{i \mid w_i > \epsilon\}$ be the set of large input items. The goal is to pack as many items from $S \cup L$ into m bins as possible. Our first observation is that if the FFI algorithm fills m bins (i.e., does not allow any more items to be packed) using only small items then it already achieves a $1 + \epsilon$ approximation.

► **Lemma 1.** *Suppose that when the FFI algorithm terminates, it has filled all bins with items of weight at most ϵ . Then FFI achieves $1 + \epsilon$ approximation ratio on this instance.*

Proof. If FFI packs all items then it clearly finds an optimal solution. Suppose that FFI rejects some items. Let w be the smallest weight of a rejected item. Thus the total remaining free space among all m bins is $< wm$ in the FFI packing. Thus, OPT can pack at most $m - 1$ more items, since it can only add items of weight $\geq w$. Let N be the number of items packed by FFI. Then we have

$$\frac{\text{OPT}}{\text{FFI}} < \frac{N + m}{N} \leq \frac{m/\epsilon + m}{m/\epsilon} = 1 + \epsilon,$$

where the second inequality follows from $N \geq m/\epsilon$, since the FFI packing uses only items of weight $\leq \epsilon$. ◀

Thus, the whole difficulty in designing a PTAS for this problem lies in the handling of large items. If FFI terminates before packing all of S , then the condition of Lemma 1 holds and hence from now on, we consider the case when FFI packs all of S . In this case an optimal solution is to pack all of S together with some subset of smallest items from L . The strategy for our algorithm is to pack a largest subset F of L that still leaves enough room to pack all of S . This means that $w(F) \leq m - w(S)$, but we also want to pack all of S efficiently. This can be guaranteed by leaving slightly more room while packing F . Namely, $w(F) \leq m(1 - \epsilon) - w(S)$ guarantees that all of S can be packed efficiently after packing F .

► **Lemma 2** (Kellerer [15]). *Suppose that we have a packing of $F \subseteq L$ such that $w(F) \leq m(1 - \epsilon) - w(S)$. Then running FFI with the packing of F as a starting point results in packing all of S .*

Proof. Initially, we have $w(F) \leq m(1 - \epsilon) - w(S) \leq m(1 - \epsilon)$. Thus, by the pigeonhole principle there is a bin with $\geq \epsilon$ free space. Thus, we can pack the first item s_1 from S . Now, we have $w(F) \leq m(1 - \epsilon) - w(S) \leq m(1 - \epsilon) - w(s_1)$, i.e., $w(F) + w(s_1) \leq m(1 - \epsilon)$. Again, by the pigeonhole principle there is a bin with $\geq \epsilon$ free space, so we can pack the second item from S , and so on. \blacktriangleleft

► **Remark.** Note that the argument in the above lemma does not use the *increasing* property of FFI. Therefore, even FF can be used to complete the partial packing F with all of S .

The next lemma shows that the extra “breathing room” that we leave to guarantee an efficient packing of S does not hurt the approximation ratio.

► **Lemma 3.** *Let F be the largest subset of L that can be packed into m bins with total weight $\leq m(1 - \epsilon) - w(S)$. Then*

$$\frac{\text{OPT}}{|F| + |S|} \leq 1 + 3\epsilon.$$

Proof. If $F = L$ then we are done. Otherwise, let $w > \epsilon$ be the smallest weight of an item from $L \setminus F$. Then $|S| \geq w(S)/\epsilon \geq w(S)/w$ and $|F| \geq \frac{m(1-\epsilon)-w(S)}{w}$. Thus, $|F| + |S| \geq \frac{m(1-\epsilon)}{w}$. The total free space after packing $F \cup S$ is $\leq \epsilon m$. Thus, OPT can pack at most $\frac{\epsilon m}{w}$ more items than $|F| + |S|$. Combining all of the above, we have

$$\frac{\text{OPT}}{|F| + |S|} \leq \frac{|F| + |S| + \epsilon m/w}{|F| + |S|} \leq \frac{m(1-\epsilon)/w + \epsilon m/w}{m(1-\epsilon)/w} = \frac{1}{1-\epsilon} \leq 1 + 3\epsilon,$$

where the last inequality holds for small ϵ ; i.e., $\epsilon \leq 2/3$. \blacktriangleleft

We shall refer to the problem of finding F as in the above lemma as the LFP (“the large F problem”).

► **Remark.** Suppose that F is an approximation to the LFP with an additive ϵm term, i.e. $|F| \geq \text{OPT}_{LFP} - \epsilon m$. Then an argument similar to the one used in the above lemma shows that F together with S still gives $1 + \Theta(\epsilon)$ approximation to the original dual bin packing problem. Thus, it suffices to find a good enough F .

Before we show how to find a good approximation to the LFP, we show how to solve the dual bin packing optimally in polynomial time when the number of distinct weights of the input items is fixed. As previously stated, this follows from the known PTAS for the makespan problem. (See section 10.2 of the Vazirani text [17].)

► **Lemma 4.** *We can solve the dual bin packing problem optimally in time $O(n^{2k}m)$ where k is the number of distinct weights of the input items.*

Proof. The algorithm is a simple dynamic programming. Let w_1, \dots, w_k be the distinct weights appearing in the input. The entire input sequence can be described by a k -tuple (n_1, \dots, n_k) , where n_i is the number of items of weight w_i and $n = \sum_i n_i$. Note that the number of different possible k -tuples with n items is $O(n^k)$. Let \mathcal{K} be the set of distinct k -tuples such that each of its element fits entirely in a single bin, i.e., (ℓ_1, \dots, ℓ_k) such that $\sum_i \ell_i w_i \leq 1$. The dynamic programming table D is going to be indexed by the number of available bins m' and a possible k -tuple (ℓ_1, \dots, ℓ_k) such that $0 \leq \ell_i \leq n_i$. The value $D[(\ell_1, \dots, \ell_k), m']$ is going to indicate the maximum number of items that can be packed from the input sequence described by the state (ℓ_1, \dots, ℓ_k) in m' bins. Let $\mathcal{L}(\ell_1, \dots, \ell_k) = \{(\ell'_1, \dots, \ell'_k) \mid \forall i \ 0 \leq \ell'_i \leq \ell_i\}$. An optimal solution to the subproblem indexed

by (ℓ_1, \dots, ℓ_k) and m' consists of a packing of some element from \mathcal{L} into a single bin, and packing the remaining input items into $m' - 1$ bins:

$$D[(\ell_1, \dots, \ell_k), m'] = \max_{(\ell'_1, \dots, \ell'_k) \in \mathcal{K} \cap \mathcal{L}(\ell_1, \dots, \ell_k)} \sum_i \ell'_i + D[(\ell_1 - \ell'_1, \dots, \ell_k - \ell'_k), m' - 1].$$

The base case is given by the states where either $m' = 0$ or $\sum_i \ell_i = 0$, in which case we cannot pack any items. The overall runtime of this algorithm is $O(n^{2k}m)$ since the dynamic programming table has $O(n^k m)$ entries and each entry can be computed in time $O(n^k)$ with appropriate preprocessing of the input data. As usual, this dynamic program can be easily modified to return the actual packing rather than the number of packed items. ◀

Let L' be the subset of the smallest items from L such that $|L'|$ is as large as possible subject to $w(L') \leq m(1 - \epsilon) - w(S)$. We would like to find F by running the dynamic programming algorithm on L' . Unfortunately, L' can have too many distinct inputs. The idea is to group items of L' into few groups depending only on ϵ , reassign all weights of elements within a single group to the weight of the largest element in that group, and run the dynamic programming algorithm on the new problem instance. Then, we will need to argue that the resulting solution is an additive ϵm approximation to the LFP.

Let $\ell = |L'|$. We can assume $\ell > m$ otherwise there is a trivial way to pack ℓ items into m bins. Assume for simplicity that $m\epsilon$ is an integer and that $k = \ell/(m\epsilon)$ is also an integer. Then, we split L' into k groups of $m\epsilon$ elements each. Let $w_{j_1} \leq w_{j_2} \leq \dots \leq w_{j_\ell}$ be the weights of elements in L' . Define L_i to be the i th group consisting of items of weights $w_{j_{1+(i-1)m\epsilon}}, \dots, w_{j_{im\epsilon}}$. Reassign the weights of elements in L_i to be $w_{j_{im\epsilon}}$. Let \tilde{w} denote the modified weights. Thus, we get an instance with k distinct weights, where $k = \ell/(m\epsilon)$. Note that $\ell \leq m/\epsilon$ since we are dealing with large items $\geq \epsilon$. We conclude that $k \leq 1/\epsilon^2$. Thus, we can solve this instance in time $O(n^{2/\epsilon^2}m)$ by Lemma 4. Let F' denote this solution. Let F denote an optimal solution to LFP with the original weights. Then, we have the following.

► **Lemma 5.** F' is feasible with respect to weights w and $|F'| \geq |F| - \epsilon m$.

Proof. Since F' is feasible with weights \tilde{w} and $\tilde{w} \geq w$, we immediately conclude that F' is feasible with respect to w . Rather than directly showing $|F'| \geq |F| - \epsilon m$, we show how to construct a set F'' from F such that $|F''| \geq |F| - \epsilon m$ and F'' is feasible with respect to \tilde{w} . This will prove the lemma, since F' is a maximum cardinality set that satisfies the feasibility constraints (i.e., $|F'| \geq |F''|$). To construct such F'' , we can simply drop all items from $F \cap L_1$ and replace all items from $F \cap L_i$ by arbitrary items from L_{i-1} for $i \geq 2$. Note that $|F''| = |F \setminus L_1| \geq |F| - \epsilon m$. Moreover, F'' is feasible with respect to \tilde{w} since we are always replacing large weight items by smaller weight items. ◀

This completes the argument that approximately solving the LFP using the reassigned weights and the dynamic programming followed by FFI on small items gives a $1 + \Theta(\epsilon)$ approximation. The running time of the dynamic programming is $O(n^{2/\epsilon^2}m)$. One can run FFI in $O(n \log n + nm)$ time. The overall running time of our PTAS algorithm is $O(n^{2/\epsilon^2}m + n \log n)$, which is clearly polynomial when ϵ is fixed. Algorithm 1 describes this PTAS.

Summarizing, in this section we proved the following theorem.

► **Theorem 6.** Algorithm 1 is a PTAS for the dual bin packing problem.

Algorithm 1 Our PTAS for the dual bin packing.

procedure DUAL BIN PACKING PTAS($w_1 \leq \dots \leq w_n, m, \epsilon$)

 Let $S = \{i \mid w_i \leq \epsilon\}$

if FFI(S, m) packs $< |S|$ items **then return** FFI(S, m)

 Let $L' = \{|S| + 1, \dots, |S| + \ell\}$ be the indices of the largest subset of L such that $w(L') \leq m(1 - \epsilon) - w(S)$

 Let $k = \ell / (m\epsilon)$

 Let \tilde{w} denote new weights where items with indices $\{|S| + (i - 1)m\epsilon + 1, \dots, |S| + im\epsilon\}$ all receive weight $w_{|S| + im\epsilon}$

 Use the algorithm of Lemma 4 to obtain a packing of items F' with modified weights \tilde{w}

 Regard F' as a packing with the original weights

 Run FF on S with the packing of F' as a starting point

return the resulting packing

4 Advice Complexity of the Online Dual Bin Packing Problem for Bounded Bit Size of Input Items

In this section, we consider the online version of the dual bin packing problem in the tape-advice model. Let s be the maximum bit-size of an input item weight. Then the input bit-length is $O(sn)$. Based on the PTAS in Algorithm 1, we develop an online algorithm that achieves $1 + \epsilon$ approximation to the dual bin packing problem with $O\left(\frac{s + \log n}{\epsilon^2}\right)$ bits of advice. Before we prove the main result of this section, we need to modify Lemma 1 to work in the online setting. Recall that Lemma 1 detects when FFI is already successful enough that we don't need to do any extra work to obtain a $1 + \epsilon$ approximation. An online algorithm does not have the ability to sort the input items, thus we would like to obtain a version of Lemma 1 that detects when FF obtains a $1 + \epsilon$ approximation. The restricted subsequence first fit (RSFF) algorithm given by Renault [16] is what we need. Let $W = w_1, \dots, w_n$ be the sequence of weights given to the online algorithm. For a value η we define W_η to be the subsequence $(w_i \mid w_i \leq \eta)$. The RSFF algorithm finds the largest value of η such that FF packs all items in W_η and then returns $FF(W_\eta)$. Without loss of generality, we may assume that η is one of the w_i .

► **Lemma 7** (Implicit in Renault [16]). *If RSFF identifies an η such that $\eta \leq \epsilon$ then RSFF achieves a $1 + \epsilon$ approximation ratio.*

By replacing FFI with RSFF in the first step of Algorithm 1, we obtain the main result of this section.

► **Theorem 8.** *There is an online algorithm achieving a $1 + \Theta(\epsilon)$ strict competitive ratio for the dual bin packing problem with $O\left(\frac{s + \log n}{\epsilon^2}\right)$ bits of advice, where s is the maximum bit-size of an input item.*

Proof. The advice is obtained by slightly modifying the PTAS from Section 3. At first, the oracle writes down the value of η identified by running RSFF on the input sequence. For later convenience, we rename η by \tilde{w}_0 . This takes $O(s)$ bits of advice. This is analogous to running FFI in the original PTAS. Recall that the PTAS creates $k \leq 1/\epsilon^2$ groups of large input items L_i for $i \in [k]$ with the corresponding rounded weights \tilde{w}_i for $i \in [k]$. The oracle appends $|L_i|$ together with \tilde{w}_i for $i \in [k]$ to the advice string. This completes the specification of the advice string. The length of the advice string is $O(s + k \log |L_i| + ks) = O\left(\frac{s + \log n}{\epsilon^2}\right)$.

It is left to see that with this advice string an online algorithm can compute a $1 + \Theta(\epsilon)$ approximate solution to the instance of the dual bin packing problem. Observe that if $\tilde{w}_0 \leq \epsilon$ then by Lemma 7 the solution obtained by running FF on all items of weight $\leq \tilde{w}_0$ achieves $1 + \epsilon$ approximation, since this gives us exactly the packing produced by RSFF. From now on, we consider the case $\tilde{w}_0 > \epsilon$. Then an optimal solution might use large items. Recall that the PTAS creates a solution to the rounded instance encoded by $(|L_1|, \dots, |L_k|)$ and weights $(\tilde{w}_1, \dots, \tilde{w}_k)$, replaces this solution with actual weights of the corresponding items and fills the rest in FF fashion with the rest of the items (see the remark immediately following Lemma 2). Thus, knowing $(|L_1|, \dots, |L_k|)$ and weights $(\tilde{w}_1, \dots, \tilde{w}_k)$ from the advice, our online algorithm can reserve place holders for items in bins according to the dynamic programming solution. We refer to this space as the preallocated space, and we refer to the complement of it as the remaining space. For example, if dynamic programming solution says that bin 1 contains ℓ_i items of weight \tilde{w}_i then the online algorithm reserves ℓ_i slots of weight \tilde{w}_i in bin 1. The preallocated space in bin 1 is $\sum_i \ell_i \tilde{w}_i$ and the remaining space in bin 1 is $1 - \sum_i \ell_i \tilde{w}_i$. Now, the algorithm is ready to process the items in the online fashion. When the algorithm receives an input item of weight $\leq \epsilon$ it packs it in the remaining space in FF fashion. When the algorithm receives an item of weight $\in (\tilde{w}_{i-1}, \tilde{w}_i]$, it packs it into the first available preallocated slot of weight \tilde{w}_i . By the construction of advice, we are guaranteed that when the algorithm is done processing the inputs, all preallocated slots are occupied and all small items are packed. By Theorem 6 this solution is a $1 + \epsilon$ approximation. \blacktriangleleft

5 Advice Complexity of the Online Dual Bin Packing Problem for General Weights

In this section we show that the online dual bin packing without any restrictions on s requires $\Theta_\epsilon(n)$ advice to approximate OPT within $1 + \epsilon$. Observe that the upper bound $O(n/\epsilon)$ immediately follows from the result of Renault [16] in the per-request advice model. A somewhat stronger upper bound, $(1 - \Omega(\epsilon))n$, follows by observing that the dual bin packing belongs to the advice complexity class AOC defined by Boyar et al. [6] and then using the results from [6]. Thus, we only need to prove that in the case of unrestricted s the lower bound of $\Omega_\epsilon(n)$ holds. For sufficiently small ϵ , we show a nearly matching lower bound of $(1 - O(\epsilon \log(1/\epsilon)))n = \Omega_\epsilon(n)$ in the tape-advice model. We establish our lower bound by providing a reduction (that preserves the precision, advice and competitive ratio) from an online problem known to require a lot of advice to the dual bin packing problem. The starting point is the binary separation problem defined by Boyar et al. [7].

► **Definition 9** (Boyar et al. [7]). The *binary separation problem* is the online problem with input $I = (n_1, y_1, \dots, y_n)$ consisting of $n = n_1 + n_2$ positive values which are revealed one by one. There is a fixed partitioning of the set of items into a subset of n_1 large items and a subset of n_2 small items, so that all large items are greater than all small items. Upon receiving an item y_i , an online algorithm for the problem must guess if y belongs to the set of small or large items. After the algorithm has made a guess, it is revealed whether the guess was correct. The goal is to maximize the number of correct guesses.

Boyar et al. [7] establish a lower bound on the advice needed to achieve competitive ratio c for the binary separation problem.

► **Theorem 10** (Boyar et al. [7]). Assume that an online algorithm solves the binary separation problem on sequences $I = (n_1, y_1, \dots, y_n)$ where the y_i are n bit numbers and does so using

at most $b(n)$ bits of advice while making at most $r(n)$ mistakes. Set $\alpha = (n - r(n))/n$. If $\alpha \in [1/2, 1)$ then $b(n) \geq (1 - H(\alpha))n$ where $H(p) = p \log(1/p) + (1 - p) \log(1/(1 - p))$.

Moreover, Boyar et al. [7] provide a reduction from the binary separation problem to the standard bin packing problem to show that achieving competitive ratio $< 9/8$ requires an online algorithm to receive $\Omega(n)$ bits of advice. A simple adaptation of this reduction allows us to derive a similar result for the dual bin packing problem. We present the details below for completeness.

► **Theorem 11.** *An online algorithm achieving a competitive ratio $1 + \epsilon$ for the dual bin packing problem with unrestricted bit size of input weights requires $(1 - O(\epsilon \log(1/\epsilon)))n = \Omega_\epsilon(n)$ bits of advice, provided $\epsilon < 1/19$.*

Proof. We show how to reduce the binary separation problem to the dual bin packing problem while preserving the size of advice and the competitive ratio.

Let ALG be an algorithm for the dual bin packing problem that achieves competitive ratio c and uses advice $b(n)$. Let $I = (n_1, (y_1, \dots, y_n))$ be an input to the binary separation problem. We define ALG' for solving I as follows. ALG' constructs an instance of the dual bin packing problem in the online fashion. It will use decisions and the advice string of ALG to make decisions about its own inputs y_i . Let $\delta_{\max} > \delta_{\min} > 0$ be small enough numbers. Suppose that we have a strictly decreasing function $f : \mathbb{R} \rightarrow (\delta_{\min}, \delta_{\max})$. ALG' invokes ALG with n bins and $2n$ items. ALG' constructs input weights to ALG in three phases.

Phase 1 (preprocessing): the first n_1 weights are defined as $1/2 + \delta_{\min}$. This is generated by ALG' prior to any inputs seen from I .

Phase 2 (online): when y_i arrives, ALG' defines a new input item to ALG of weight $1/2 - f(y_i)$. In this phase ALG' uses decisions of ALG to handle its own inputs. If ALG packs the current item into a bin that contains $1/2 + \delta_{\min}$ item from phase 1 then ALG' declares y_i to be large. Otherwise, ALG' declares the item to be small. We shall refer to $1/2 - f(y_i)$ weight items corresponding to truly small (large) y_i as small items (large items).

Phase 3 (post processing): once ALG' has processed the entire sequence I , it appends weights $1/2 + f(y_i)$ for all truly small y_i from I . We refer to these weights as the complementary weights of small items.

First observe that OPT for the constructed instance of the dual bin packing packs all $2n$ items into n bins: the n_1 weights corresponding to the large items can be paired up with the n_1 items from phase 1, and the n_2 weights corresponding to the small items can be paired up with their complementary weights from phase 3 in the remaining $n_2 = n - n_1$ bins.

Clearly, the advice complexity and the precision of the input items are preserved by this reduction. Thus, to finish the argument we need to analyze how many mistakes ALG' does. We bound the number of mistakes in terms of the number of items unpacked by ALG. We define the following variables.

- Let p_1 be the number of items from phase 1 that were not packed by ALG.
- Let ℓ_2 be the number of large items from phase 2 that were not packed by ALG.
- Let s_2 be the number of small items from phase 2 that were not packed by ALG.
- Let p_3 be the number of items from phase 3 that were not packed by ALG.

The overall number of items that were not packed by ALG is $p_1 + \ell_2 + s_2 + p_3 \leq \frac{c-1}{c}2n$. Observe that the complementary weights can only be paired up with the corresponding small item weights, and phase 1 items can only be paired up with large or small phase 2 items.

The number of bins containing phase 1 items is $n_1 - p_1$. The number of bins containing phase 3 items is $n_2 - p_3$. Due to the above observations, all these bins have to be distinct.

Thus, the number of bins that do not contain either phase 1 or phase 3 items is $p_1 + p_3$. Call these the leftover bins. There are two types of mistakes that ALG' can do: (1) it classifies a large item as being small, and (2) it classifies a small item as being large. Since large items can only be paired either with phase 1 items or be placed in the leftover bins, type (1) mistakes occur only when large items are placed in the leftover bins or when large items remain unpacked. There can be at most $2(p_1 + p_3)$ large items in the leftover bins. Thus, ALG' makes at least $g_1 := n_1 - p_1 - 2(p_1 + p_3) - \ell_2$ correct guesses for large items. A type (2) mistake happens only when a small item is paired up with a phase 1 item. Since there can be at most $n_1 - p_1 - g_1 = 2(p_1 + p_3) + \ell_2$ phase 1 items not paired up with large items, there can be at most that many type (2) mistakes. Thus, ALG' makes at least $g_2 = n_2 - s_2 - 2(p_1 + p_3) - \ell_2$ correct guesses for small items. Overall, ALG' makes $g_1 + g_2 = n_1 + n_2 - s_2 - p_1 - 4(p_1 + p_3) - 2\ell_2 \geq n_1 + n_2 - 5(p_1 + \ell_2 + s_2 + p_3) \geq n - 10\frac{c-1}{c}n$ good guesses. The fraction of good guesses is then $\frac{n-10(c-1)n/c}{n} = \frac{10-9c}{c}$. By Theorem 10, it follows that $b(n) \geq (1 - H((10-9c)/c))n$. Observe that $(10-9c)/c \in (1/2, 1)$ provided that $c \in (1, 20/19)$. In particular, if ϵ is a small positive constant, then achieving a competitive ratio $c = 1 + \epsilon$ for the dual bin packing problem requires $(1 - H((1-9\epsilon)/(1+\epsilon)))n = (1 - H(O(\epsilon)))n = (1 - O(\epsilon \log(1/\epsilon)))n = \Omega_\epsilon(n)$ bits of advice. \blacktriangleleft

All in all, the dual bin packing problem admits short advice in case of s bounded by a slowly growing function of n , but requires long advice when s is unrestricted. This is akin to the distinction between the polynomial time vs pseudo-polynomial time in the regular Turing machine world. One of the conceptual contributions of this paper is a demonstration that “pseudo-short” advice and “truly short” advice are provably different. To make this idea precise, we introduce two natural classes of efficient advice problems.

► **Definition 12.** The class *EAC* (*efficient advice complexity*) consists of online problems P such that an input to P is given by n items, and the advice complexity of achieving $1 + \epsilon$ competitive ratio for P is $O_\epsilon(\text{poly}(\log n))$.

Denoting the maximum bit size of an input item to P by s , we define a superclass *WEAC* (*weakly efficient advice complexity*) of *EAC* to consist of those online problems P such that the advice complexity of achieving $1 + \epsilon$ competitive ratio for P is $O_\epsilon(\text{poly}(\log n, s))$.

EAC class is defined by analogy with communication complexity where $O(\text{poly}(\log n))$ communication is considered efficient (see Babai et al. [2]). *WEAC* class is also natural. The advice length bound of algorithms for *WEAC* problems suggests that the advice can consist of a short description of combinatorial parameters of a problem (e.g., length of a stream, index into a stream, which take $O(\log n)$ bits to describe) plus a small (polylogarithmic) number of actual data items from the stream.

In light of the above definitions and the main result of this section and Section 4, the dual bin packing problem witnesses the following class separation theorem.

► **Theorem 13.** $WEAC \neq EAC$.

6 Conclusion

We presented a simple PTAS for the dual bin packing problem. Although a PTAS for a more general multiple knapsack problem was already known, our PTAS is arguably simpler to state and analyze. Its simplicity helped us to adapt it to the tape-advice model of online algorithms. We showed that a $1 + \epsilon$ competitive ratio for the dual bin packing problem is achievable with $O\left(\frac{s + \log n}{\epsilon^2}\right)$ bits of tape advice. We showed that the dependence on s

is necessary to obtain such small advice, as the dual bin packing problem requires $\Omega_\epsilon(n)$ when $\epsilon > 0$ is small enough, s is unrestricted, and m is part of the input. We introduced two natural advice complexity classes EAC and WEAC. The conceptual distinction between the classes WEAC and EAC is similar to the Turing machine world distinction between pseudo-polynomial time and strongly polynomial time. EAC captures problems that can be approximated to within $1 + \epsilon$ with $O_\epsilon(\text{poly} \log n)$ bits of advice, whereas WEAC captures problems that can be approximated to within $1 + \epsilon$ with $O_\epsilon(\text{poly}(\log n, s))$ bits of advice. Our results on the dual bin packing problem imply that $\text{WEAC} \neq \text{EAC}$.

One immediate question left open by our work is whether there is a small advice algorithm for small s which requires less advice bits. More specifically, does there exist a $1 + \epsilon$ approximation using $o_\epsilon(s) + O_\epsilon(\log n)$ advice bits for $s = o(n)$? In this paper we exclusively studied the dual bin packing in the regime of obtaining $1 + \epsilon$ competitive ratio when ϵ is small and m is part of the input. Are there sublinear advice algorithms for large ϵ , e.g., $\epsilon = 1/2$? Also, does the dual bin packing admit sublinear advice algorithms when m is a small constant? It is also interesting to see whether or not results for the dual bin packing problem can be extended to more general problems such as when bins have different capacities, and more generally to the multiple knapsack problem, while preserving conceptual simplicity. Last and perhaps a most important question is whether or not there exist online algorithms with efficiently computable (i.e., linear or even online computable as in [12, 4]) advice for the dual bin packing problem achieving a constant competitive ratio.

Acknowledgements. We thank the anonymous reviewers, especially reviewer 2, for their helpful comments.

References

- 1 S.F Assmann, D.S Johnson, D.J Kleitman, and J.Y.-T Leung. On a dual version of the one-dimensional bin packing problem. *J. of Algorithms*, 5(4):502 – 525, 1984. URL: <http://www.sciencedirect.com/science/article/pii/019667748490004X>, doi:[http://dx.doi.org/10.1016/0196-6774\(84\)90004-X](http://dx.doi.org/10.1016/0196-6774(84)90004-X).
- 2 Laszlo Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Proc. of the 27th Symp. on Found. of Comput. Sci.*, SFCS '86, pages 337–347, 1986.
- 3 Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Kráľovič, Richard Kráľovič, and Tobias Mömke. On the advice complexity of online problems. *Algorithms and Computation*, pages 331–340, 2009.
- 4 Allan Borodin, Denis Pankratov, and Amirali Salehi-Abari. On conceptually simple algorithms for variants of online bipartite matching. In *WAOA'17: The 15th workshop on approximation and online algorithms (To appear)*, 2017.
- 5 Joan Boyar, Lene M Favrrholdt, Christian Kudahl, Kim S Larsen, and Jesper W Mikkelsen. Online algorithms with advice: a survey. *ACM SIGACT News*, 47(3):93–129, 2016.
- 6 Joan Boyar, Lene M. Favrrholdt, Christian Kudahl, and Jesper W. Mikkelsen. The advice complexity of a class of hard online problems. *Theory of Comput. Sys.*, 2016.
- 7 Joan Boyar, Shahin Kamali, Kim S. Larsen, and Alejandro López-Ortiz. Online bin packing with advice. *Algorithmica*, 74(1):507–527, Jan 2016.
- 8 Joan Boyar, Kim S. Larsen, and Morten N. Nielsen. The accommodating function: A generalization of the competitive ratio. *SIAM J. on Comput.*, 31(1):233–258, 2001.
- 9 Chandra Chekuri and Sanjeev Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM J. on Comput.*, 35(3):713–728, 2005.

- 10 J. Csirik and V. Totik. Online algorithms for a dual version of bin packing. *Discrete Applied Mathematics*, 21(2):163 – 167, 1988. URL: <http://www.sciencedirect.com/science/article/pii/0166218X88900522>, doi:[http://dx.doi.org/10.1016/0166-218X\(88\)90052-2](http://dx.doi.org/10.1016/0166-218X(88)90052-2).
- 11 Marek Cygan, Lukasz Jez, and Jiri Sgall. Online knapsack revisited. *Theory Comput. Sys.*, 58(1):153–190, 2016.
- 12 Christoph Dürr, Christian Konrad, and Marc Renault. On the Power of Advice and Randomization for Online Bipartite Matching. In *Proc. of ESA*, pages 37:1–37:16, 2016.
- 13 Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. *Theoretical Computer Science*, 412(24):2642–2656, 2011.
- 14 Edward G. Coffman Jr., Joseph Y.-T. Leung, and D. W. Ting. Bin packing: Maximizing the number of pieces packed. *Acta Inf.*, 9:263–271, 1978.
- 15 Hans Kellerer. A polynomial time approximation scheme for the multiple knapsack problem. In *Proc. of RANDOM-APPROX*, volume 1671, pages 51–62. Springer, 1999.
- 16 Marc P Renault. Online algorithms with advice for the dual bin packing problem. *Central Eur. J. of Op. Res.*, pages 1–14, 2016.
- 17 Vijay V. Vazirani. *Approximation algorithms*. Springer, 2001.